

Regularized Matrix Regression

Hua Zhou

Department of Statistics
North Carolina State University
Raleigh, NC 27695-8203
hau_zhou@ncsu.edu

Lexin Li

Department of Statistics
North Carolina State University
Raleigh, NC 27695-8203
li@stat.ncsu.edu

Abstract

Modern technologies are producing a wealth of data with complex structures. For instance, in two-dimensional digital imaging, flow cytometry, and electroencephalography, matrix type covariates frequently arise when measurements are obtained for each combination of two underlying variables. To address scientific questions arising from those data, new regression methods that take matrices as covariates are needed, and sparsity or other forms of regularization are crucial due to the ultrahigh dimensionality and complex structure of the matrix data. The popular lasso and related regularization methods hinge upon the sparsity of the true signal in terms of the number of its nonzero coefficients. However, for the matrix data, the true signal is often of, or can be well approximated by, a low rank structure. As such, the sparsity is frequently in the form of low rank of the matrix parameters, which may seriously violate the assumption of the classical lasso. In this article, we propose a class of regularized matrix regression methods based on spectral regularization. Highly efficient and scalable estimation algorithm is developed, and a degrees of freedom formula is derived to facilitate model selection along the regularization path. Superior performance of the proposed method is demonstrated on both synthetic and real examples.

Key Words: Electroencephalography; multidimensional array; Nesterov method; nuclear norm; spectral regularization; tensor regression.

1 Introduction

Modern scientific applications are frequently producing data sets where the sampling unit is not in the form of a vector but instead a matrix. Examples include two-dimensional digital imaging data, which contain the quantized brightness value of a given color at a number of rows and columns of pixels; and flow cytometric data, which consist of the fluorescence intensity of multiple cells at multiple channels. Our motivating example is a study of an electroencephalography (EEG) data of alcoholism (<http://kdd.ics.uci.edu/datasets/eeg/eeg.data.html>).

The study consists of 122 subjects with two groups: an alcoholic group and a normal control group, and each subject was exposed to a stimulus. Voltage values were measured from 64 channels of electrodes placed on the subject’s scalp for 256 time points, so each sampling unit is a 256×64 matrix. It is of scientific interest to study the association between alcoholism and the pattern of voltage over times and channels (Li et al., 2010). The generalized linear model (GLM) (McCullagh and Nelder, 1983) offers a useful tool for that purpose, where the response Y is the binary indicator of alcoholic or control, and the predictors include the matrix-valued EEG data \mathbf{X} and possible covariate vector \mathbf{Z} such as age and gender. However, the classical GLM deals with a vector of covariates, and the presence of matrix type covariates poses fresh challenges to statistical analysis. First, naïvely turning a matrix into a vector results in an exceedingly large dimensionality; for instance, for the EEG data, the dimension is $p = 256 \times 64 = 16,384$, whereas the sample size is only $n = 122$. Second, vectorization destroys the wealth of structural information inherently possessed in the matrix data; e.g., it is commonly expected that the voltage values of the adjacent time points and channels are highly correlated. Given the ultrahigh dimensionality and the complex structure, *regularization* becomes crucial for the analysis of such data. In this article, we propose a novel regularization solution for regression with matrix covariates that efficiently tackles the ultrahigh dimensionality while preserving the matrix structure.

A large variety of regularization methods have been developed in recent years. Among them, penalization has been playing a powerful role in stabilizing the estimates, improving the risk property, and increasing the generalization power in classical regressions. Popular penalization techniques include lasso (Tibshirani, 1996; Donoho and Johnstone, 1994), SCAD (Fan and Li, 2001), fused lasso (Tibshirani et al., 2005), elastic net (Zou and Hastie, 2005), and many others. For regressions with matrix covariates, a direct approach is to first vectorize the covariates then apply the classical penalization techniques. Regularization helps alleviate the problem that the dimensionality far exceeds the sample size. However, this is unsatisfactory, since it does not incorporate the matrix structural information. More importantly, the solution is based upon a fundamental assumption that the true underlying signal is sparse in terms of the ℓ_0 norm of the regression parameters. In matrix regressions, however, it is often the case that the true signal is of a low rank structure, or can be well approximated by a low rank structure. As such, sparsity is in terms of the *rank* of the matrix parameters, which is intrinsically different from sparsity in the number of nonzero entries.

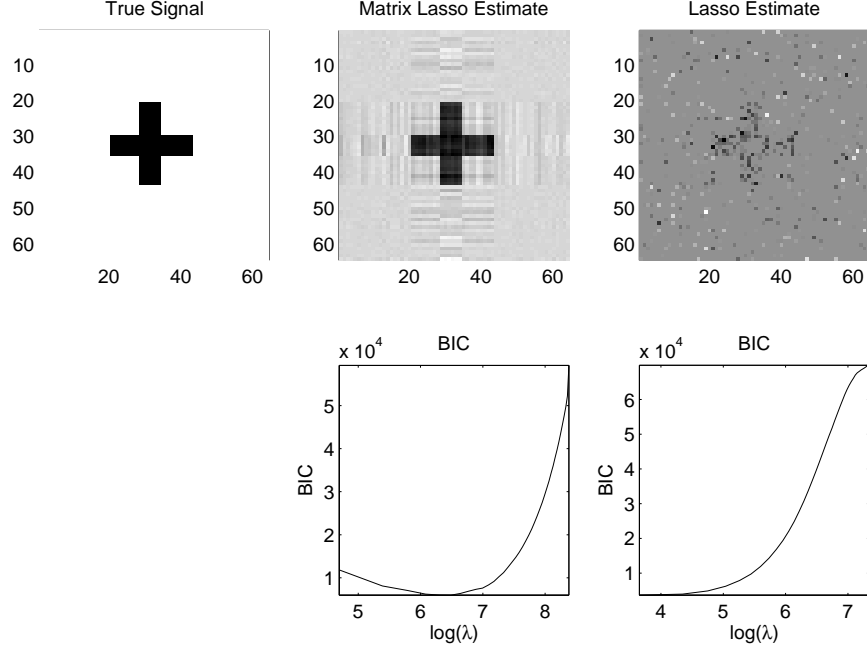


Figure 1: Comparison of the nuclear norm regularized estimation to the classical lasso. The matrix covariate is 64×64 , and the sample size is 500. Top panel: True signal (left), nuclear norm regularized estimate with minimal BIC (center), and the classical lasso estimate with minimal BIC (right). Bottom panel: BIC along solution paths for the nuclear norm regularization (left) and the classical lasso regularization (right).

To see how such a difference affects signal estimation in matrix regressions, we consider the following illustrative example. We generated a normal response Y with mean, $\mu = \gamma^\top \mathbf{Z} + \langle \mathbf{B}, \mathbf{X} \rangle$, and variance one. $\mathbf{Z} \in \mathbb{R}^5$ denotes a usual vector of covariates with standard normal entries, and $\gamma = (1, \dots, 1)^\top$. $\mathbf{X} \in \mathbb{R}^{64 \times 64}$ denotes the matrix covariates, of which all entries are standard normal, and \mathbf{B} is the coefficient matrix of the same size. \mathbf{B} is binary, with the true signal region, which is a cross shape in our example, equal to one and the rest zero. The inner product between two matrices is defined as $\langle \mathbf{B}, \mathbf{X} \rangle = \langle \text{vec} \mathbf{B}, \text{vec} \mathbf{X} \rangle = \sum_{i_1, i_2} \beta_{i_1 i_2} x_{i_1 i_2}$, where $\text{vec}(\cdot)$ is the vectorization operator that stacks the columns of a matrix into a vector. We sampled $n = 500$ instances $\{(y_i, \mathbf{x}_i, \mathbf{z}_i), i = 1, \dots, 500\}$, and our goal is to identify \mathbf{B} through a regression of y_i on $(\mathbf{x}_i, \mathbf{z}_i)$. We note that our problem differs from the usual edge detection or object recognition in imaging processing (Qiu, 2005, 2007). In our setup, all elements of the image \mathbf{X} follow the same distribution. The signal region is defined through the coefficient image \mathbf{B} and needs to be inferred from the association between Y and \mathbf{X} after adjusting for \mathbf{Z} . We also note that the total number of entries in \mathbf{B} is $4096 = 64^2$, and the number of nonzero ones is 240 (about 5.8%). We applied two approaches. The first is

lasso to the vectorized \mathbf{X} . That is, we solve the optimization problem

$$\min_{\mathbf{B}} \frac{1}{2} \sum_{i=1}^n (y_i - \gamma^\top \mathbf{z}_i - \langle \mathbf{B}, \mathbf{x}_i \rangle)^2 + \lambda \|\text{vec} \mathbf{B}\|_1,$$

where $\|\text{vec} \mathbf{B}\|_1$ is the ℓ_1 norm of the vectorized \mathbf{B} , and λ is the regularization parameter. The lower right panel of Figure 1 displays the Bayesian information criterion (BIC) along the lasso solution path, which suggests a model with the maximum number of predictors (500) that is allowed given the sample size. The parameter estimate under this model is shown in the upper right panel, which appears far away from the truth. The second solution we consider is penalizing the nuclear norm of \mathbf{B} , i.e., we solve

$$\min_{\mathbf{B}} \frac{1}{2} \sum_{i=1}^n (y_i - \gamma^\top \mathbf{z}_i - \langle \mathbf{B}, \mathbf{x}_i \rangle)^2 + \lambda \|\mathbf{B}\|_*,$$

where the nuclear norm $\|\mathbf{B}\|_* = \sum_j \sigma_j(\mathbf{B})$, and $\sigma_j(\mathbf{B})$'s are the singular values of the matrix \mathbf{B} . The nuclear norm $\|\mathbf{B}\|_*$ is a suitable measure of the “size” of a matrix parameter, and is a convex relaxation of $\text{rank}(\mathbf{B}) = \|\sigma(\mathbf{B})\|_0$. This is analogous to the ℓ_1 norm for a vector (Recht et al., 2010). The lower left graph of Figure 1 displays the BIC along the solution path of the nuclear norm penalized matrix regression, and the upper middle panel shows the corresponding estimate with minimal BIC. It is clearly seen that the nuclear norm estimate achieves a substantially better recovery than the lasso estimate. One might argue that fused lasso (Tibshirani et al., 2005) might give a better recovery of such piecewise constant signals. However, there are numerous low rank signals, e.g., $(01 \dots 01)^\top (10 \dots 10)$, which are extremely non-smooth and would fail fused lasso.

More generally, in this article, we propose a family regularized regression models with matrix covariates based on spectral regularization. Our contributions are multifold. First, we employ a spectral regularization formulation and integrate within a generalized linear model (GLM) framework. The resulting model works for a variety of penalization functions, including lasso, elastic net, SCAD, and many others, as well as different types of response variables, including normal, binary and count outcomes. Second, we develop a highly efficient and scalable Nesterov algorithm for model estimation with explicit, non-asymptotic convergence rate. We emphasize that such a highly scalable algorithm is critical for analyzing large-scale and ultrahigh-dimensional matrix data. Third, we derive the effective degrees of freedom of selected models, which is crucial for tuning of the regularization parameter. The result can be viewed as an extension of the degrees of freedom development from the

classical lasso model (Zou et al., 2007) and the group lasso model (Yuan and Lin, 2006) to the generalized linear matrix model. On the other hand, our proposal is *not* simply another variant of lasso and alike. We aim at *matrix regression* problems, which are important in imaging and other scientific applications but have received relatively little attention.

Our proposal is related to but also distinct from two recent developments involving matrix data. The first is a recent proposal of a family of generalized linear models with matrix or tensor (multidimensional arrays) covariates (Zhou et al., 2012). The basic idea is to impose a particular low rank structure (CP decomposition) on \mathbf{B} and then introduce a sparse penalty on the coefficients of \mathbf{B} . That solution fits the model at a *fixed* rank of the matrix/tensor regression parameters, and thus corresponds to the hard thresholding in the classical vector covariate case. In contrast, our solution does not fix the rank of the parameters and is a soft thresholding procedure. Moreover, even when using a convex penalty function such as the lasso penalty, the approach of Zhou et al. (2012) involves a challenging non-convex optimization task, whereas the solution in this paper remains a convex problem. We also note that Hung and Wang (2011) considered matrix logistic regression, which is a special case of Zhou et al. (2012), and they did not investigate any sparsity regularization. The second related work is the line of research in matrix completion, where nuclear norm type regularization has been widely employed (Candès and Recht, 2009; Mazumder et al., 2010; Cai et al., 2010). However, the two approaches are different in that, the matrix completion problem aims to recover a low rank matrix when only a small portion of its entries are observed., whereas our approach concerns about regressions with matrix covariates.

The rest of the article is organized as follows. We formulate the spectral regularization for matrix regression in Section 2, and develop a highly scalable algorithm for the associated optimization in Section 3. We derive the degrees of freedom formula in Section 4, and investigate the numerical performance of the proposed method in Section 5. We conclude the paper with a discussion of potential future research in Section 6. We delegate all the technical proofs to the Appendix.

2 Spectral Regularization

We first fix the notations. For any matrix $\mathbf{B} \in \mathbb{R}^{p_1 \times p_2}$, $\sigma(\mathbf{B}) = (\sigma_1(\mathbf{B}), \dots, \sigma_q(\mathbf{B}))$, $q = \min\{p_1, p_2\}$, denotes the vector of decreasingly ordered singular value mapping of \mathbf{B} . That is $\sigma_1(\mathbf{B}) \geq \sigma_2(\mathbf{B}) \geq \dots \geq \sigma_r(\mathbf{B}) > \sigma_{r+1}(\mathbf{B}) = \dots = \sigma_q(\mathbf{B}) = 0$, where $r = \text{rank}(\mathbf{B})$. Let

Y denote the response variable, $\mathbf{Z} \in \mathbb{R}^{p_0}$ the vector covariate, $\mathbf{X} \in \mathbb{R}^{p_1 \times p_2}$ the 2D matrix covariate, and $(y, \mathbf{x}, \mathbf{z})$ their sample instances.

We consider the generalized linear model setup, where Y belongs to an exponential family with probability mass function or density

$$p(y|\mathbf{x}, \mathbf{z}) = \exp \left\{ \frac{y\theta - b(\theta)}{a(\phi)} + c(y, \phi) \right\},$$

and the first conditional moment is $E(Y|\mathbf{X}, \mathbf{Z}) = \mu = b'(\theta)$, and μ is of the form

$$g(\mu) = \eta = \boldsymbol{\gamma}^\top \mathbf{Z} + \langle \mathbf{B}, \mathbf{X} \rangle, \quad (1)$$

where g is a known link function, $\boldsymbol{\gamma} \in \mathbb{R}^{p_0}$, and $\mathbf{B} \in \mathbb{R}^{p_1 \times p_2}$. For simplicity, in our subsequent development of the regularized matrix model estimation, we drop the vector covariate \mathbf{Z} and its associated parameter $\boldsymbol{\gamma}$. However, the results can be extended straightforwardly to incorporate \mathbf{Z} and $\boldsymbol{\gamma}$. Also, we only consider GLM with a univariate response, whereas extensions to more complex models such as quasi-likelihood models and multivariate responses are straightforward.

For generality in terms of penalty, we consider the spectral regularization problem

$$\min_{\mathbf{B}} h(\mathbf{B}) = \ell(\mathbf{B}) + J(\mathbf{B}), \quad (2)$$

where $\ell(\mathbf{B})$ is a loss function; for the GLM, we use the negative log-likelihood as the loss. $J(\mathbf{B}) = f \circ \sigma(\mathbf{B})$, where $f : \mathbb{R}^q \rightarrow \mathbb{R}$ is a function of the singular values of \mathbf{B} . The choice $f(\mathbf{w}) = \lambda \sum_{j=1}^q |w_j|$ and the least squares loss corresponds to the special case of the nuclear norm regularization problem we considered in the Introduction. In general, for sparsity of the spectrum, f takes the general form

$$f(\mathbf{w}) = \sum_{j=1}^q P_\eta(|w_j|, \lambda),$$

where P is a scalar penalty function, η is the parameter indexing the penalty family, and λ is the tuning constant. We list some commonly used penalty functions below.

- Power family (Frank and Friedman, 1993)

$$P_\eta(|w|, \lambda) = \lambda |w|^\eta, \quad \eta \in (0, 2].$$

Two important special cases of this family are the lasso penalty when $\eta = 1$ (Tibshirani, 1996; Chen et al., 2001) and the ridge penalty when $\eta = 2$ (Hoerl and Kennard, 1970).

- Elastic net (Zou and Hastie, 2005)

$$P_\eta(|w|, \lambda) = \lambda \{ (\eta - 1)w^2/2 + (2 - \eta)|w| \}, \quad \eta \in [1, 2].$$

Varying η from 1 to 2 bridges the lasso to the ridge penalty.

- Log penalty (Candès et al., 2008; Armagan et al., 2011)

$$P_\eta(|w|, \lambda) = \lambda \ln(\eta + |w|), \quad \eta > 0.$$

- SCAD (Fan and Li, 2001), in which the penalty is defined via its partial derivative

$$\frac{\partial}{\partial |w|} P_\eta(|w|, \lambda) = \lambda \left\{ 1_{\{|w| \leq \lambda\}} + \frac{(\eta\lambda - |w|)_+}{(\eta - 1)\lambda} 1_{\{|w| > \lambda\}} \right\}, \quad \eta > 2.$$

Integration shows SCAD as a natural quadratic spline with knots at λ and $\eta\lambda$

$$P_\eta(|w|, \lambda) = \begin{cases} \lambda|w| & |w| < \lambda \\ \lambda^2 + \frac{\eta\lambda(|w| - \lambda)}{\eta - 1} - \frac{w^2 - \lambda^2}{2(\eta - 1)} & |w| \in [\lambda, \eta\lambda] \\ \lambda^2(\eta + 1)/2 & |w| > \eta\lambda \end{cases}.$$

For small signals $|w| < \lambda$, it acts as lasso; for large signals $|w| > \eta\lambda$, the penalty flattens and leads to the unbiasedness of the regularized estimate.

- MC+ penalty (Zhang, 2010), which is similar to SCAD and is defined by the partial derivative

$$\frac{\partial}{\partial |w|} P_\eta(|w|, \lambda) = \lambda \left(1 - \frac{|w|}{\lambda\eta} \right)_+.$$

Integration shows that the penalty function

$$P_\eta(|w|, \lambda) = \left(\lambda|w| - \frac{w^2}{2\eta} \right) 1_{\{|w| < \lambda\eta\}} + \frac{\lambda^2\eta}{2} 1_{\{|w| \geq \lambda\eta\}}, \quad \eta > 0,$$

is quadratic on $[0, \lambda\eta]$ and flattens beyond $\lambda\eta$. Varying η from 0 to ∞ bridges hard thresholding (ℓ_0 regression) to lasso (ℓ_1) shrinkage.

We also comment that, besides the above sparsity penalties, other forms of regularization can be useful, depending on the scientific question of interest. For instance, the choice $f(\mathbf{w}) = \lambda \sum_{j=1}^{q-1} |w_j - w_{j+1}| = \lambda(w_1 - w_r)$ produces the regularization for the “spiked” matrix model, i.e., matrices with clustered eigen-/singular values (Johnstone, 2001).

Convexity is essential for studying convergence properties of optimization problems. We first state the necessary and sufficient condition for the convexity of the regularizer J . Its proof follows from the theory of spectral function (Borwein and Lewis, 2006) and is given in the Appendix.

Lemma 1. *The functional $J(\mathbf{B}) = f \circ \sigma(\mathbf{B})$ is convex and lower semicontinuous if and only if f is convex and lower semicontinuous. Furthermore, for a convex f , the subdifferential of J at \mathbf{B} , which admits singular value decomposition $\mathbf{U} \text{diag}(\mathbf{b}) \mathbf{V}^\top$, is*

$$\partial J(\mathbf{B}) = \partial(f \circ \sigma)(\mathbf{B}) = \mathbf{U} \text{diag}[\partial f(\mathbf{b})] \mathbf{V}^\top.$$

Lemma 1 immediately leads to the optimality condition when both loss and regularizer are convex.

Theorem 1. *When both the loss ℓ and f are convex, all local minima of the regularized program (2) are global minimum and is unique if ℓ is strictly convex. A matrix $\mathbf{B} = \mathbf{U} \text{diag}(\mathbf{b}) \mathbf{V}^\top$ is a global minimum if and only if*

$$\mathbf{0}_{p_1 \times p_2} \in \nabla \ell(\mathbf{B}) + \mathbf{U} \text{diag}[\partial f(\mathbf{b})] \mathbf{V}^\top.$$

When either the loss ℓ or f is non-convex, the regularized objective function (2) may be non-convex and there lacks an easy-to-check optimality condition.

3 Estimation Algorithm

We utilize the powerful Nesterov optimal gradient method (Nesterov, 1983, 2004) for minimizing the non-smooth and possibly non-convex objective function (2). We first state a matrix thresholding formula for spectral regularization, which forms the building blocks of the Nesterov algorithm.

Proposition 1. *For a given matrix \mathbf{A} with singular value decomposition $\mathbf{A} = \mathbf{U} \text{diag}(\mathbf{a}) \mathbf{V}^\top$, the optimal solution to*

$$\min_{\mathbf{B}} \frac{1}{2} \|\mathbf{B} - \mathbf{A}\|_{\text{F}}^2 + f \circ \sigma(\mathbf{B})$$

shares the same singular vectors as \mathbf{A} and its ordered singular values are the solution to

$$\min_{\mathbf{b}} \frac{1}{2} \|\mathbf{b} - \mathbf{a}\|_2^2 + f(\mathbf{b}).$$

An immediate consequence of Proposition 1 is the following well-known singular value thresholding formula for nuclear norm regularization (Cai et al., 2010).

Corollary 1. *For a given matrix \mathbf{A} with singular value decomposition $\mathbf{A} = \mathbf{U}\text{diag}(\mathbf{a})\mathbf{V}^\top$, the optimal solution to*

$$\min_{\mathbf{B}} \frac{1}{2} \|\mathbf{B} - \mathbf{A}\|_{\text{F}}^2 + \lambda \|\mathbf{B}\|_*$$

shares the same singular vectors as \mathbf{A} and its singular values are $b_i = (a_i - \lambda)_+$.

Given the above matrix thresholding formula, we are ready to present the Nesterov algorithm for minimizing (2). The Nesterov method has attracted increasing attention in recent years due to its efficiency in solving regularization problems (Beck and Teboulle, 2009b). It resembles the classical gradient descent algorithm in that only the first order gradients of the objective function are utilized to produce next algorithmic iterate from current search point, and as such is simple to implement. It differs from the gradient descent algorithm by extrapolating the previous two algorithmic iterates to generate the next search point. This extrapolation step incurs trivial computational cost but improves the convergence rate dramatically. It has been shown to be optimal among a wide class of convex smooth optimization problems (Nemirovski, 1994; Nesterov, 2004).

We summarize our Nesterov method for solving spectral regularization problem (2) in Algorithm 1. Each iteration consists of three steps: (i) predict a search point \mathbf{S} by a linear extrapolation from previous two iterates (line 3 of Algorithm 1), (ii) perform gradient descent from the search point \mathbf{S} possibly with Armijo type line search (lines 4-10), and (iii) force the descent property of the next iterate (lines 11-15).

In step (i), $\alpha^{(t)}$ is a scalar sequence that plays a critical role in the extrapolation. We update this sequence as in the original Nesterov method (line 16 of Algorithm 1), whereas other sequences, for instance $\alpha^{(t)} = (t-1)/(t+2)$, can also be used. In step (ii), the gradient descent is based on the first order approximation to the loss function at the current search point $\mathbf{S}^{(t)}$,

$$\begin{aligned} g(\mathbf{B}|\mathbf{S}^{(t)}, \delta) &= \ell(\mathbf{S}^{(t)}) + \langle \nabla \ell(\mathbf{S}^{(t)}), \mathbf{B} - \mathbf{S}^{(t)} \rangle + \frac{1}{2\delta} \|\mathbf{B} - \mathbf{S}^{(t)}\|_{\text{F}}^2 + J(\mathbf{B}) \\ &= \frac{1}{2\delta} \|\mathbf{B} - [\mathbf{S}^{(t)} - \delta \nabla \ell(\mathbf{S}^{(t)})]\|_{\text{F}}^2 + J(\mathbf{B}) + c^{(t)}, \end{aligned}$$

where the constant δ is determined during the line search and the constant $c^{(t)}$ collects terms irrelevant to the optimization. The “ridge” term $(2\delta)^{-1} \|\mathbf{B} - \mathbf{S}^{(t)}\|_{\text{F}}^2$ acts as a trust region and shrinks the next iterate towards $\mathbf{S}^{(t)}$. If the loss function $\ell \in \mathcal{C}_{1,1}$, which denotes the class of functions that are convex, continuously differentiable and the gradient satisfies

```

1 Initialize  $\mathbf{B}^{(0)} = \mathbf{B}^{(1)}$ ,  $\delta > 0$ ,  $\alpha^{(0)} = 0$ ,  $\alpha^{(1)} = 1$  ;
2 repeat
3    $\mathbf{S}^{(t)} \leftarrow \mathbf{B}^{(t)} + \left( \frac{\alpha^{(t-1)} - 1}{\alpha^{(t)}} \right) (\mathbf{B}^{(t)} - \mathbf{B}^{(t-1)})$  ;
4   repeat
5      $\mathbf{A}_{\text{temp}} \leftarrow \mathbf{B}^{(t)} - \delta \nabla \ell(\mathbf{S}^{(t)})$  ;
6     Compute SVD  $\mathbf{A}_{\text{temp}} = \mathbf{U} \text{diag}(\mathbf{a}) \mathbf{V}^\top$  ;
7      $\mathbf{b} \leftarrow \text{argmin}_{\mathbf{x}} (2\delta)^{-1} \|\mathbf{x} - \mathbf{a}\|_2^2 + f(\mathbf{x})$  ;
8      $\mathbf{B}_{\text{temp}} \leftarrow \mathbf{U} \text{diag}(\mathbf{b}) \mathbf{V}^\top$  ;
9      $\delta \leftarrow \delta/2$  ;
10  until  $h(\mathbf{B}_{\text{temp}}) \leq g(\mathbf{B}_{\text{temp}} | \mathbf{S}^{(t)}, \delta)$  ;
11  if  $h(\mathbf{B}_{\text{temp}}) \leq h(\mathbf{B}^{(t)})$  then
12     $\mathbf{B}^{(t+1)} \leftarrow \mathbf{B}_{\text{temp}}$  ;
13  else
14     $\mathbf{B}^{(t+1)} \leftarrow \mathbf{B}^{(t)}$  ;
15  end
16   $\alpha^{(t+1)} \leftarrow (1 + \sqrt{1 + (2\alpha^{(t)})^2})/2$ 
17 until objective value converges;

```

Algorithm 1: Nesterov method for spectral regularized matrix regression (2).

$\|\nabla \ell(\mathbf{u}) - \nabla \ell(\mathbf{v})\| \leq \mathcal{L}(\ell) \|\mathbf{u} - \mathbf{v}\|$ with a known gradient Lipschitz constant $\mathcal{L}(\ell)$ for all \mathbf{u}, \mathbf{v} , then δ is fixed at $\mathcal{L}(\ell)^{-1}$. In practice, the gradient Lipschitz constant is often unknown. Then δ is updated dynamically to capture the unknown $\mathcal{L}(\ell)$ using the classical Armijo line search rule (Nocedal and Wright, 2006; Lange, 2004). Solution to the surrogate function $g(\mathbf{B} | \mathbf{S}^{(t)}, \delta)$ is obtained by Proposition 1. Singular value decomposition is performed on the intermediate matrix $\mathbf{A}_{\text{temp}} = \mathbf{B}^{(t)} - \delta \nabla \ell(\mathbf{S}^{(t)})$. The next iterate $\mathbf{B}^{(t+1)}$ shares the same singular vectors as \mathbf{A} and its singular values $\mathbf{b}^{(t+1)}$ are determined by minimizing $(2\delta)^{-1} \|\mathbf{b} - \mathbf{a}\|_2^2 + f(\mathbf{b})$, where $\mathbf{a} = \sigma(\mathbf{A}_{\text{temp}})$. For a nuclear norm regularization $f(\mathbf{w}) = \lambda \sum_j |w_i|$, the solution is given by soft thresholding the singular values $b_i^{(t+1)} = (a_i - \lambda\delta)_+$. In this special case, only the top singular values/vectors need to be retrieved. The Lanczos method (Golub and Van Loan, 1996) is extremely efficient for this purpose. For a linear regularization function $f(\mathbf{w}) = \lambda \|\mathbf{D}\mathbf{w}\|_1$ where \mathbf{D} has full column rank, reparameterization $\mathbf{c} = \mathbf{D}\mathbf{b}$ turns the problem to: $\min_{\mathbf{c}} \frac{1}{2\delta} \|(\mathbf{D}^\top \mathbf{D})^{-1} \mathbf{D}^\top \mathbf{c} - \mathbf{a}\|_2^2 + \lambda \|\mathbf{c}\|_1$, which is a standard lasso problem with many efficient solvers available. When \mathbf{D} does not have a full column rank, we append extra rows such that the expanded matrix, denoted by $\tilde{\mathbf{D}}$, has full column rank and then solve the above lasso problem with $\tilde{\mathbf{D}}$ and only part of \mathbf{c} penalized.

For minimization of a smooth convex function ℓ in $\mathcal{C}_{1,1}$, it is well-known that the Nesterov

method is optimal with the convergence rate at order $O(t^{-2})$, where t indicates the iterate number. In contrast, the gradient descent has a slower convergence rate of $O(t^{-1})$. Our nuclear norm regularization problem (2) is non-smooth, but the same convergence result can be established, which is summarized in Theorem 2. Its proof is omitted for brevity, while readers are referred to Beck and Teboulle (2009b).

Theorem 2. *Suppose ℓ is continuously differentiable with a gradient Lipschitz constant $\mathcal{L}(\ell)$. Let $\mathbf{B}^{(t)}$ be the iterates generated by the Nesterov method described in Algorithm 1. Then the objective value $h(\mathbf{B}^{(t)})$ monotonically converges. Furthermore, if J is convex, then*

$$h(\mathbf{B}^{(t)}) - h(\mathbf{B}^*) \leq \frac{4\mathcal{L}(\ell)\|\mathbf{B}^{(0)} - \mathbf{B}^*\|_{\text{F}}^2}{(t+1)^2} \quad (3)$$

for all $t \geq 0$ and any minimum point \mathbf{B}^* .

We make a few important remarks here. The first remark regards the monotonicity of the objective function during iterations. Because of the extrapolation step, the objective values of algorithmic iterates $f(\mathbf{B}^{(t)})$ are not guaranteed to be monotonically decreasing. When the loss $\ell \in \mathcal{C}_{1,1}$ and the regularizer J is convex, convergence of the objective values is guaranteed with the explicit convergence rate (3). Because of potential use of a non-convex J , we enforce monotonicity of algorithmic iterates (lines 11-15 in Algorithm 1), which is essential for the convergence of at least the objective values. After each gradient descent step, if the new iterate fails to decrease the objective value, then the current iterate is same as the previous one. In other words, the next gradient descent is initiated from the previous iterate. Fortunately, the fast convergence rate (3) still holds under the assumptions $\ell \in \mathcal{C}_{1,1}$ and J is convex. See Beck and Teboulle (2009a) for the argument.

The second remark is about the non-convex loss function. The Nesterov method and its convergence properties hinge upon convexity of the loss ℓ . It covers many commonly used statistical models, including linear model and GLMs with canonical links. For GLM with non-canonical link, the loss function may be non-convex, which could cause trouble in the Nesterov method. The iteratively reweighted least squares strategy can be applied in this scenario. At each IWLS step, the Nesterov method is used to solve the penalized weighted least squares problem, which is convex.

The final remark is about an efficient way for estimating the Lipschitz constant L for the GLM loss. Each step halving in the line search part of Algorithm 1 involves an expensive singular value decomposition. Therefore even a rough initial estimate of L potentially cuts

the computational cost significantly. Recall that a twice differentiable function f is L -Lipschitz continuously differentiable if and only if $\mathbf{v}^\top d^2 f(\mathbf{u}) \mathbf{v} \leq L \|\mathbf{v}\|_2^2$ for all \mathbf{v} . The Fisher information matrix of a GLM model with systematic part (1) is: $\mathbf{I}(\mathbf{B}) = \mathbb{E}[d^2 \ell(\mathbf{B})] = \sum_{i=1}^n \omega_i (\text{vec} \mathbf{x}_i)(\text{vec} \mathbf{x}_i)^\top$, where $\omega_i = \{\mu'_i(\eta_i)/\sigma_i\}^2$, η_i is the systematic part, μ_i is the mean, and σ_i^2 is the variance corresponding to the i th observation. Then in light of the Cauchy-Schwartz inequality

$$\mathbf{v}^\top \mathbf{I}(\mathbf{B}) \mathbf{v} = \sum_i \omega_i (\mathbf{v}^\top \text{vec} \mathbf{x}_i)(\mathbf{v}^\top \text{vec} \mathbf{x}_i)^\top \leq \sum_i \omega_i \|\text{vec} \mathbf{x}_i\|_2^2 \|\mathbf{v}\|_2^2 = \|\mathbf{v}\|_2^2 \left(\sum_i \omega_i \|\mathbf{x}_i\|_F^2 \right),$$

and thus an initial estimate of L is given by $L \approx \sum_i \{\mu'_i(\eta_i)\}^2 / \sigma_i^2 \|\mathbf{X}_i\|_F^2$.

4 Degrees of Freedom

In this section, we address the problem of choosing the tuning parameter λ that yields the best model along the regularization path according to certain criteria. Cross validation is commonly used for parameter tuning in practice. However, for large data, it may incur considerable computation burden. There exist computationally attractive alternatives, such as Akaike information criterion (AIC) (Akaike, 1974) and Bayesian information criterion (BIC) (Schwarz, 1978), which often yield performance comparable to cross validation in practice.

Consider a normal model under the GLM (1). For simplicity, we again drop the covariate vector \mathbf{Z} :

$$Y = \langle \mathbf{X}, \mathbf{B} \rangle + \epsilon \quad (4)$$

where ϵ is a normal error with mean zero and variance σ^2 . Let y_i denote the i th observation of Y , and $\hat{y}_i(\lambda)$ denote the estimated response under a given tuning parameter λ from the minimization of (2). Then for this normal model, AIC and BIC are defined by

$$\begin{aligned} \text{AIC}(\lambda) &= \frac{\sum_i \{y_i - \hat{y}_i(\lambda)\}^2}{\sigma^2} + 2\text{df}(\lambda) \\ \text{BIC}(\lambda) &= \frac{\sum_i \{y_i - \hat{y}_i(\lambda)\}^2}{\sigma^2} + \ln(n)\text{df}(\lambda). \end{aligned}$$

In applications, the variance σ^2 is often unknown but can be estimated from the fitted value by least squares estimation. An essential element in the above model selection criteria is the

effective degrees of freedom $\text{df}(\lambda)$ of the selected model. Using Stein's theory of unbiased risk estimation (Stein, 1981), Efron (2004) showed that

$$\text{df}(\lambda) = \mathbb{E} \left\{ \text{tr} \left(\frac{\partial \hat{\mathbf{y}}}{\partial \mathbf{y}} \right) \right\} = \frac{1}{\sigma^2} \sum_{i=1}^n \text{cov}(\hat{y}_i(\lambda), y_i)$$

with expectation taken with respect to Y , $\mathbf{y} = (y_1, \dots, y_n)^\top$, and $\hat{\mathbf{y}} = (\hat{y}_1(\lambda), \dots, \hat{y}_n(\lambda))^\top$. This formulation has been productively used to derive the degrees of freedom estimate in least angle regression (Efron et al., 2004), lasso (Zou et al., 2007), group penalized regression (Yuan and Lin, 2006), and sign-coherent group penalized regression (Chiquet et al., 2011).

We derive a degrees of freedom estimate for the nuclear norm regularized estimate under normal model. For an orthonormal design, i.e., $\mathbf{\Xi}^\top \mathbf{\Xi} = \mathbf{I}_{p_1 p_2}$ with the matrix $\mathbf{\Xi}$ having rows $(\text{vec} \mathbf{x}_i)^\top$, the derived estimate is unbiased for the true degrees of freedom. In practice, it yields results comparable to cross validation even for non-orthogonal designs. The technical proof is relegated to the Appendix.

Theorem 3. *Assume that the data is generated from model (4) with $\text{vec} \mathbf{x}_i$ orthonormal. Consider the nuclear norm regularized estimate*

$$\hat{\mathbf{B}}_\lambda = \underset{\mathbf{B}}{\text{argmin}} \frac{1}{2} \sum_i (y_i - \langle \mathbf{X}_i, \mathbf{B} \rangle)^2 + \lambda \|\mathbf{B}\|_*$$

with singular values $\sigma(\hat{\mathbf{B}}_\lambda) = (b_1(\lambda), \dots, b_q(\lambda))$ where $q = \min\{p_1, p_2\}$. Let $\hat{\mathbf{B}}_{\text{LS}}$ be the usual least squares estimate and assume that it has distinct positive singular values $\sigma_1 > \dots > \sigma_q > 0$. With the convention $\sigma_i = 0$ for $i > q$, the following expression is an unbiased estimate of the degree of freedom of the regularized fit

$$\hat{\text{df}}(\lambda) = \sum_{i=1}^q 1_{\{b_i(\lambda) > 0\}} \left(1 + \sum_{1 \leq j \leq p_1, j \neq i} \frac{\sigma_i(\sigma_i - \lambda)}{\sigma_i^2 - \sigma_j^2} + \sum_{1 \leq j \leq p_2, j \neq i} \frac{\sigma_i(\sigma_i - \lambda)}{\sigma_i^2 - \sigma_j^2} \right).$$

This formula for the degrees of freedom is interesting in several aspects. First it does not involve any information on the singular vectors of the least squares estimate, but only requires singular values. Second, $\hat{\text{df}}(\lambda)$ is continuous in λ , in contrast to the piecewise constant degrees of freedom estimate for the classical lasso (Zou et al., 2007). Third, at $\lambda = 0$, $\hat{\mathbf{B}}_0 = \hat{\mathbf{B}}_{\text{LS}}$ almost surely has a full rank and the degrees of freedom is

$$\begin{aligned} & \sum_{i=1}^q \left(1 + \sum_{1 \leq j \leq p_1, j \neq i} \frac{\sigma_i^2}{\sigma_i^2 - \sigma_j^2} + \sum_{1 \leq j \leq p_2, j \neq i} \frac{\sigma_i^2}{\sigma_i^2 - \sigma_j^2} \right) \\ &= q + q(q-1) + q(p_1 + p_2 - 2q) = p_1 p_2, \end{aligned}$$




Figure 2: Degrees of freedom estimate $\widehat{\text{df}}(\lambda)$ versus the number of parameters in the estimated model $\widehat{\mathbf{B}}_\lambda$ with a $\widehat{\mathbf{B}}_{\text{LS}} \in \mathbb{R}^{64 \times 64}$.

which is exactly the number of parameters without any regularization.

Figure 2 plots the estimated degrees of freedom $\widehat{\text{df}}(\lambda)$ for a $\widehat{\mathbf{B}}_{\text{LS}} \in \mathbb{R}^{64 \times 64}$, as well as the naïve count of the number of free parameters in the fitted matrix parameter $\widehat{\mathbf{B}}_\lambda$ of rank $r(\lambda)$, which equals $r(\lambda)(p_1 + p_2) - r^2(\lambda)$. The estimated degrees of freedom appears smaller than the naïve count, reflecting the overwhelming shrinkage effect over model searching. At $\lambda = 0$, it coincides with the number of matrix elements $64^2 = 4096$, reflecting the effect of no shrinkage.

Finally we notice that the degrees of freedom estimate in Theorem 3 is limiting as it requires existence of the least squares estimate $\widehat{\mathbf{B}}_{\text{LS}}$ which is not true when $n < p_1 p_2$. In this case we may use a ridge estimate $\widehat{\mathbf{B}}_{\text{ridge}(\tau)}$ where

$$\text{vec} \widehat{\mathbf{B}}_{\text{ridge}(\tau)} = (\mathbf{\Xi}^\top \mathbf{\Xi} + \tau \mathbf{I}_{p_1 p_2})^{-1} \mathbf{\Xi}^\top \mathbf{y},$$

which always exists and is unique. Assume that $\widehat{\mathbf{B}}_{\text{ridge}(\tau)}$ admits a singular value decomposition $\widehat{\mathbf{B}}_{\text{ridge}(\tau)} = \mathbf{U} \text{diag}(\boldsymbol{\sigma}) \mathbf{V}^\top$. The following degree of freedom formula

$$\begin{aligned} \widehat{\text{df}}(\tau) = \sum_{i=1}^q 1_{\{b_i(\tau) > 0\}} & \left[1 + \frac{1}{1 + \tau} \sum_{1 \leq j \leq p_1, j \neq i} \frac{\sigma_i \{(1 + \tau) \sigma_i - \lambda\}}{\sigma_i^2 - \sigma_j^2} \right. \\ & \left. + \frac{1}{1 + \tau} \sum_{1 \leq j \leq p_2, j \neq i} \frac{\sigma_i \{(1 + \tau) \sigma_i - \lambda\}}{\sigma_i^2 - \sigma_j^2} \right] \end{aligned}$$

generalizes Theorem 3 and is unbiased for the true degree of freedom under the same assumptions as Theorem 3. Its proof is given in the Appendix.

5 Numerical Examples

We have conducted intensive numerical studies with two aims: first, we investigate the empirical performance of the proposed spectral regularized regression with matrix covariates, and second, we compare with the corresponding classical regularization solutions. Four methods are under comparison: a matrix regression with nuclear norm regularization (since it takes the form $f(\mathbf{w}) = \lambda \sum_{j=1}^q |w_j|$ in the regularization problem (2), we call this solution matrix lasso, or simply, m.lasso), a usual vector regression after vectorizing the matrix covariate with a lasso penalty (lasso), a matrix regression with power spectral regularization (matrix power, or simply, m.power), and the corresponding vector regression with a power penalty (power). For the power penalty, we fixed the coefficient $\eta = 0.5$. Our goal here is not to best tune η . Instead, we examine this penalty since it yields a nearly unbiased estimate. We also examined another unbiased penalty, SCAD, which yields very similar results as power $\eta = 0.5$, and thus its results are not reported here for brevity. We also note that the lasso penalty is a convex penalty, while the power $\eta = 0.5$ is non-convex. We summarize our findings in three examples. First, we elaborated on the illustrative example by examining a number of different geometric and natural shapes. Second, we generated some synthetic data and compared different regularization solutions under varying ranks and sparsity. Lastly, we revisited the motivating electroencephalography (EEG) data analysis mentioned in the Introduction.

Example 1: 2D Shapes. We elaborated on the illustrative example in the Introduction, by employing the same model setup, but examining a variety of signal shapes. We presented the true signal followed by the estimates from the four aforementioned regularization methods, where the regularization parameter λ was tuned by BIC. The signals of square, cross, T-shape are given in Figure 3, where the true signal is of a low rank structure (rank 1 or 2). The signals of triangle, disk and butterfly are given in Figure 4, where the true signal is not of an exact low rank structure, however, can be well approximated by so (Zhou et al., 2012). It is seen that the matrix version of regularized estimators clearly outperform their vector version counterparts, for both lasso and power penalties. Comparing matrix lasso with matrix power, the two yield comparable results, while the former is better for the high rank signals, and the latter is better for the low rank signals. We will further verify this observation in the next simulation example.

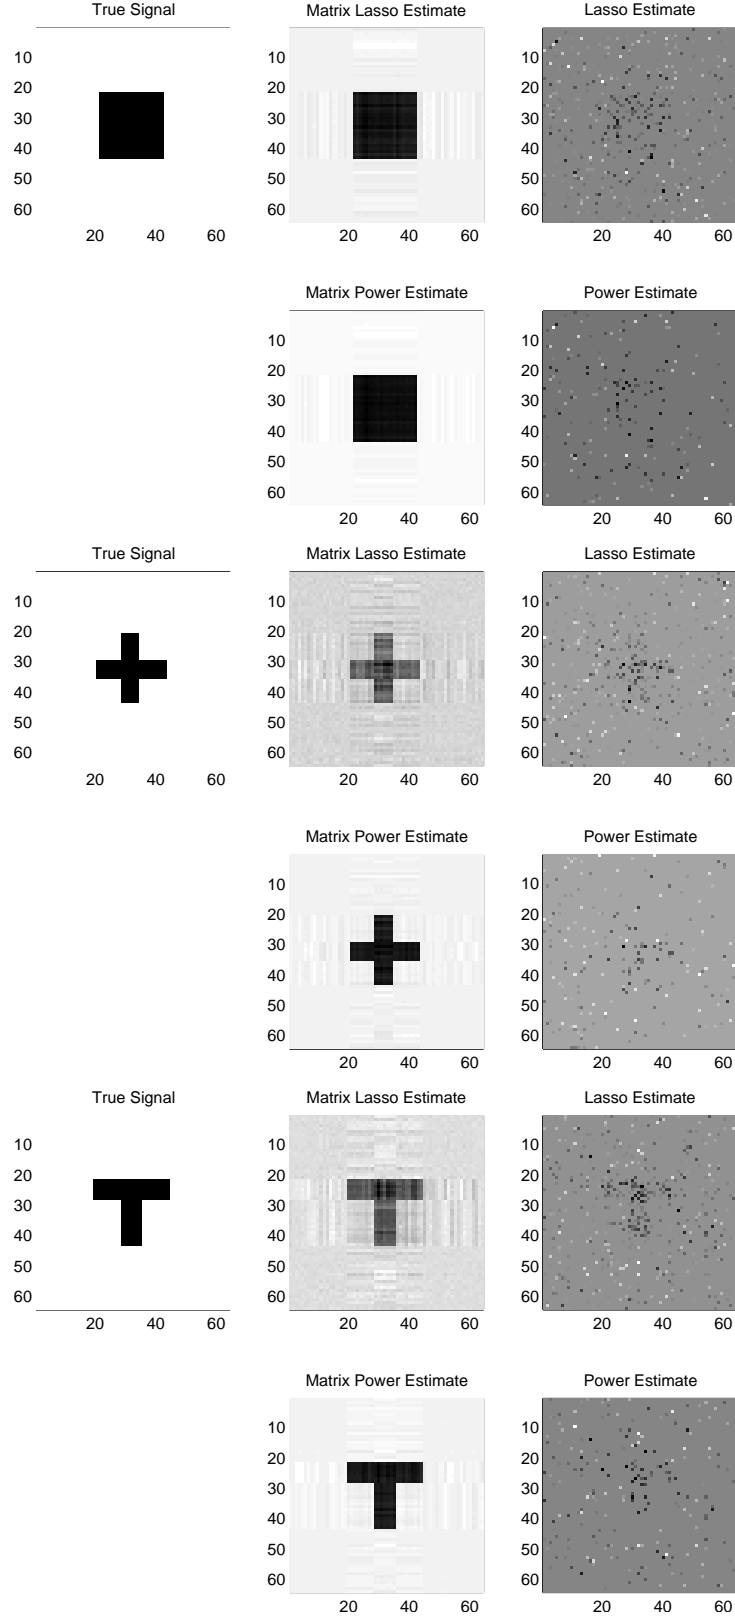


Figure 3: Comparison of the matrix and vector version of regularized estimators for low rank signals.

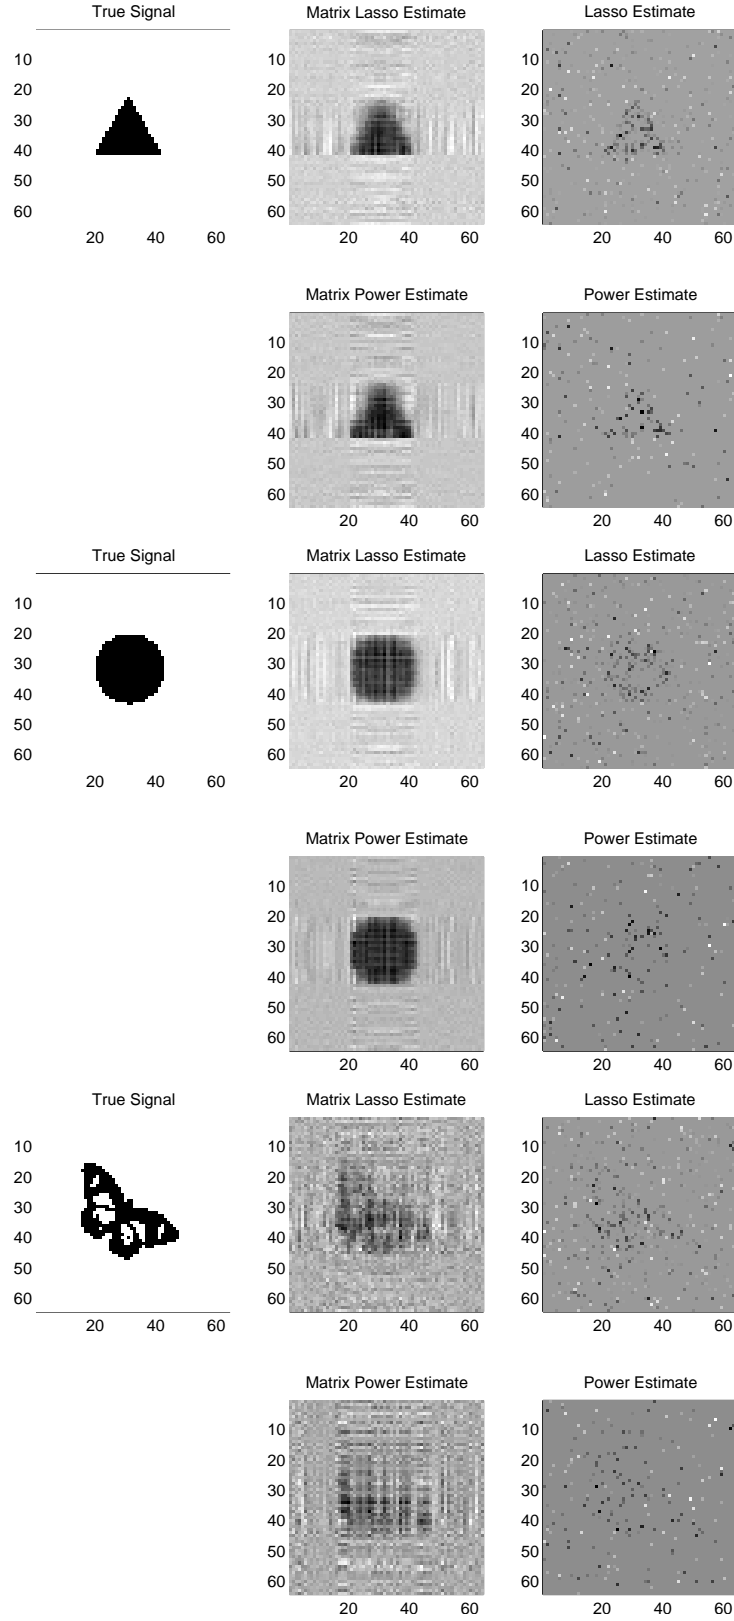


Figure 4: Comparison of the matrix and vector version of regularized estimators for high rank signals.

Table 1: Parameter estimation of a normal model. Reported are the mean and standard deviation (in the parenthesis) of the RMSE for \mathbf{B} out of 100 data replications.

Sparsity		Rank			
s	Method	$R = 1$	$R = 5$	$R = 10$	$R = 20$
1%	m.lasso	0.031 (0.006)	0.074 (0.016)	0.086 (0.015)	0.090 (0.013)
	lasso	0.022 (0.017)	0.022 (0.010)	0.023 (0.013)	0.022 (0.006)
	m.power	0.010 (0.001)	0.063 (0.021)	0.097 (0.021)	0.107 (0.017)
	power	0.009 (0.025)	0.005 (0.001)	0.007 (0.017)	0.006 (0.001)
5%	m.lasso	0.044 (0.007)	0.172 (0.021)	0.199 (0.022)	0.212 (0.020)
	lasso	0.212 (0.049)	0.214 (0.035)	0.215 (0.032)	0.217 (0.026)
	m.power	0.010 (0.001)	0.171 (0.029)	0.236 (0.029)	0.254 (0.024)
	power	0.270 (0.060)	0.271 (0.042)	0.272 (0.039)	0.275 (0.032)
10%	m.lasso	0.051 (0.008)	0.250 (0.028)	0.288 (0.023)	0.309 (0.023)
	lasso	0.320 (0.046)	0.343 (0.040)	0.343 (0.034)	0.345 (0.030)
	m.power	0.009 (0.001)	0.258 (0.041)	0.351 (0.033)	0.375 (0.032)
	power	0.399 (0.056)	0.427 (0.050)	0.428 (0.044)	0.426 (0.036)
20%	m.lasso	0.062 (0.011)	0.350 (0.031)	0.418 (0.031)	0.450 (0.031)
	lasso	0.475 (0.050)	0.523 (0.050)	0.539 (0.049)	0.534 (0.044)
	m.power	0.009 (0.001)	0.354 (0.041)	0.511 (0.039)	0.562 (0.044)
	power	0.585 (0.062)	0.645 (0.060)	0.667 (0.060)	0.663 (0.055)
50%	m.lasso	0.087 (0.015)	0.556 (0.030)	0.700 (0.044)	0.792 (0.040)
	lasso	0.775 (0.044)	1.054 (0.081)	1.112 (0.089)	1.130 (0.069)
	m.power	0.010 (0.001)	0.493 (0.042)	0.760 (0.055)	0.903 (0.062)
	power	0.952 (0.057)	1.300 (0.107)	1.365 (0.113)	1.388 (0.082)

Example 2: Synthetic Data. We consider a class of synthetic signals to compare various regularization methods under different ranks and sparsity levels. Specifically, we generated the matrix covariates \mathbf{X} of size 64×64 and the 5-dimensional vector covariates \mathbf{Z} , both of which consist of independent standard normal entries. We set the sample size $n = 500$, whereas the number of parameters is $64 \times 64 + 5 = 4,101$. We set $\boldsymbol{\gamma} = (1, \dots, 1)^\top$, and generated the true array signal as $\mathbf{B} = \mathbf{B}_1 \mathbf{B}_2^\top$, where $\mathbf{B}_d \in \mathbb{R}^{p_d \times R}$, $d = 1, 2$. R controls the rank of the generated signal. Moreover, each entry of \mathbf{B} is 0/1, and the percentage of non-zero entries is controlled by a sparsity level constant s . That is, each entry of \mathbf{B}_d is a Bernoulli with probability of one equal to $\sqrt{1 - (1 - s)^{1/R}}$. We varied the rank $R = 1, 5, 10$ and 20, and the level of (non)sparsity $s = 0.01, 0.05, 0.1, 0.2$ and 0.5. (So $s = 0.05$ means about 5% of entries of \mathbf{B} are ones and the rest are zeros.) We generated both a normal and a binomial response Y with the systematic part as in (1), $\mu = \eta$ for the normal model, and

Table 2: Prediction of a normal model. Reported are the mean and standard deviation (in the parenthesis) of the RMSE for y out of 100 data replications.

Sparsity		Rank			
s	Method	$R = 1$	$R = 5$	$R = 10$	$R = 20$
1%	m.lasso	1.707 (0.195)	4.597 (1.091)	5.525 (1.009)	5.865 (0.826)
	lasso	1.610 (0.978)	1.544 (0.487)	1.622 (0.762)	1.530 (0.233)
	m.power	1.175 (0.042)	3.996 (1.164)	5.534 (1.047)	5.992 (0.874)
	power	1.236 (1.240)	1.133 (0.742)	1.168 (0.962)	1.066 (0.039)
5%	m.lasso	2.278 (0.309)	11.031 (1.418)	12.837 (1.578)	13.754 (1.352)
	lasso	13.488 (2.923)	13.625 (2.115)	13.688 (2.047)	13.855 (1.609)
	m.power	1.177 (0.043)	10.253 (1.479)	13.013 (1.593)	14.118 (1.407)
	power	14.713 (3.384)	14.877 (2.109)	14.933 (2.018)	15.113 (1.597)
10%	m.lasso	2.542 (0.420)	16.101 (1.940)	18.590 (1.568)	19.810 (1.514)
	lasso	19.601 (2.661)	21.257 (2.488)	21.165 (2.052)	21.187 (1.742)
	m.power	1.179 (0.045)	15.386 (2.161)	18.840 (1.588)	20.239 (1.597)
	power	21.138 (2.879)	22.796 (2.782)	22.708 (2.323)	22.546 (1.776)
20%	m.lasso	3.188 (0.655)	22.723 (2.157)	26.783 (2.264)	28.874 (2.256)
	lasso	28.687 (3.017)	31.862 (3.219)	32.588 (3.182)	32.449 (2.767)
	m.power	1.169 (0.042)	21.401 (2.309)	27.103 (2.171)	29.271 (2.189)
	power	30.879 (3.284)	34.350 (3.495)	35.013 (3.600)	34.998 (2.994)
50%	m.lasso	4.566 (1.026)	35.651 (2.153)	45.132 (2.989)	50.916 (3.054)
	lasso	45.815 (2.749)	62.834 (4.896)	66.490 (5.523)	67.478 (4.444)
	m.power	1.180 (0.047)	29.638 (2.079)	42.104 (2.381)	48.892 (2.945)
	power	49.798 (3.091)	68.073 (5.719)	71.626 (6.063)	72.729 (4.986)

$\mu = 1/\{1 + \exp(-\eta)\}$ for the binomial.

We evaluated the performance of each method from two aspects: parameter estimation and prediction. For the former, we employed BIC for parameter tuning and the root mean squared error (RMSE) as the evaluation criterion. For the latter, we used an independent validation data set to tune the regularization parameter and a testing data set to evaluate the prediction error, which is measured by the RMSE of the response for the normal case, and the mis-classification error rate for the binomial case. Those are all common practices in the literature. We replicated the experiment 100 times, and report the mean and standard deviation of the corresponding criterion out of 100 replications in Tables 1 – 4. The best outcomes are highlighted in bold face. Since the RMSE for the vector-valued coefficient γ shows the same qualitative pattern as that for the array coefficient \mathbf{B} , we only present the results for \mathbf{B} for brevity.

Table 3: Parameter estimation of a binomial model. Reported are the mean and standard deviation (in the parenthesis) of the RMSE for \mathbf{B} out of 100 data replications.

Sparsity		Rank			
s	Method	$R = 1$	$R = 5$	$R = 10$	$R = 20$
1%	m.lasso	0.087 (0.023)	0.094 (0.021)	0.096 (0.015)	0.094 (0.014)
	lasso	0.095 (0.024)	0.097 (0.022)	0.097 (0.016)	0.095 (0.014)
	m.power	0.077 (0.020)	0.091 (0.020)	0.094 (0.015)	0.093 (0.013)
	power	0.095 (0.024)	0.097 (0.022)	0.098 (0.016)	0.095 (0.014)
5%	m.lasso	0.210 (0.042)	0.222 (0.032)	0.231 (0.027)	0.231 (0.021)
	lasso	0.220 (0.042)	0.225 (0.032)	0.233 (0.027)	0.233 (0.021)
	m.power	0.193 (0.042)	0.216 (0.032)	0.226 (0.027)	0.228 (0.021)
	power	0.221 (0.042)	0.225 (0.032)	0.233 (0.027)	0.233 (0.021)
10%	m.lasso	0.304 (0.043)	0.333 (0.035)	0.339 (0.031)	0.331 (0.028)
	lasso	0.315 (0.043)	0.336 (0.035)	0.342 (0.031)	0.333 (0.028)
	m.power	0.287 (0.043)	0.326 (0.035)	0.334 (0.031)	0.327 (0.028)
	power	0.315 (0.043)	0.337 (0.035)	0.342 (0.031)	0.333 (0.028)
20%	m.lasso	0.438 (0.044)	0.502 (0.047)	0.510 (0.038)	0.515 (0.039)
	lasso	0.449 (0.044)	0.506 (0.047)	0.513 (0.038)	0.517 (0.039)
	m.power	0.419 (0.044)	0.494 (0.047)	0.504 (0.038)	0.510 (0.038)
	power	0.449 (0.044)	0.507 (0.047)	0.513 (0.038)	0.518 (0.039)
50%	m.lasso	0.695 (0.038)	0.990 (0.072)	1.044 (0.076)	1.056 (0.057)
	lasso	0.706 (0.038)	0.997 (0.073)	1.049 (0.077)	1.061 (0.058)
	m.power	0.676 (0.039)	0.979 (0.072)	1.034 (0.076)	1.048 (0.057)
	power	0.706 (0.038)	0.997 (0.073)	1.049 (0.077)	1.061 (0.058)

We make the following observations. First, for the normal response, the proposed matrix version of estimators almost always outperform the corresponding vector version in terms of both parameter estimation and prediction, and this holds true for both the lasso and power penalties. The only exception is that when the signal is extremely sparse ($s = 1\%$ of non-zero elements), where the usual vector version of regularized estimators perform slightly better. Second, for the binomial response, the matrix version is superior than the vector version for all ranks and sparsity levels, and such a superiority is more striking in terms of predicting the binary outcome. Third, as for the effects of the signal rank and sparsity, the regular version of regularized estimators perform better when the signal is more sparse (a smaller s), while is relatively insensitive to the rank (R). By contrast, the proposed matrix version of regularized estimators perform better when the rank is smaller, and is insensitive to the coefficient sparsity. Those patterns agree with our expectations since the former penalizes

Table 4: Prediction of a binomial model. Reported are the mean and standard deviation (in the parenthesis) of the misclassification error for y out of 100 data replications.

Sparsity		Rank			
s	Method	$R = 1$	$R = 5$	$R = 10$	$R = 20$
1%	m.lasso	0.231 (0.029)	0.337 (0.030)	0.362 (0.024)	0.372 (0.024)
	lasso	0.411 (0.038)	0.416 (0.037)	0.413 (0.035)	0.415 (0.031)
	m.power	0.252 (0.028)	0.346 (0.027)	0.370 (0.025)	0.382 (0.029)
	power	0.763 (0.075)	0.764 (0.071)	0.777 (0.082)	0.771 (0.068)
5%	m.lasso	0.199 (0.021)	0.363 (0.023)	0.393 (0.022)	0.409 (0.026)
	lasso	0.459 (0.022)	0.460 (0.025)	0.460 (0.030)	0.461 (0.023)
	m.power	0.220 (0.025)	0.368 (0.023)	0.398 (0.024)	0.416 (0.025)
	power	0.853 (0.091)	0.854 (0.088)	0.824 (0.080)	0.850 (0.080)
10%	m.lasso	0.194 (0.023)	0.359 (0.027)	0.396 (0.024)	0.404 (0.027)
	lasso	0.469 (0.022)	0.466 (0.022)	0.467 (0.023)	0.462 (0.020)
	m.power	0.217 (0.027)	0.368 (0.027)	0.404 (0.024)	0.413 (0.024)
	power	0.864 (0.079)	0.865 (0.095)	0.860 (0.086)	0.874 (0.076)
20%	m.lasso	0.193 (0.030)	0.343 (0.027)	0.375 (0.027)	0.396 (0.026)
	lasso	0.469 (0.022)	0.466 (0.025)	0.470 (0.025)	0.465 (0.023)
	m.power	0.217 (0.030)	0.355 (0.024)	0.384 (0.027)	0.403 (0.023)
	power	0.886 (0.094)	0.870 (0.091)	0.859 (0.089)	0.858 (0.089)
50%	m.lasso	0.187 (0.023)	0.282 (0.033)	0.312 (0.032)	0.344 (0.027)
	lasso	0.472 (0.019)	0.471 (0.025)	0.470 (0.023)	0.465 (0.023)
	m.power	0.214 (0.023)	0.308 (0.030)	0.336 (0.028)	0.359 (0.024)
	power	0.872 (0.091)	0.860 (0.089)	0.869 (0.083)	0.873 (0.092)

directly on the coefficient sparsity, whereas the latter penalizes on the rank sparsity. Finally, comparing the lasso penalty with the power penalty, the two yield similar results, while the lasso usually performs better when the rank is large, and the power penalty is better when the rank is small. Such patterns agree with what we observed in Example 1, and can provide some useful guidelines when choosing a penalty function given the data.

Example 3: EEG Data Analysis. We analyzed the motivating EEG data in the Introduction and compared a number of regularization estimators. The data set consists of 77 individuals with alcoholism and 44 individuals as the control. For each subject, 64 channels of electrodes were placed at different locations of scalp and the voltage values are recorded at 256 time points. Besides, each subject performed 120 trials under three types of stimuli: single stimulus, two matched stimuli and two unmatched stimuli. One primary interest was to study the association between alcoholism and the pattern of voltage values over times and




Figure 5: EEG signals of an alcoholic subject (left) and a control subject (right). Each curve tracks signals from one electrode.

channels. Figure 5 displays the EEG signals from a randomly chosen alcoholic subject and a control subject, where the horizontal axis is the time, the vertical axis is the voltage, and each curve represents an electrode channel. It is clearly seen that the alcoholic and control subjects demonstrate distinct characteristics.

Li et al. (2010) and Hung and Wang (2011) both analyzed this same data set. Following their practice, we focused on the data under the single stimulus condition only, and averaged for all the trials under that condition for each subject. The resulting covariates \mathbf{x}_i are 256×64 matrices, and the response y_i is a binary variable indicating whether the i th subject is alcoholic ($y_i = 1$) or not ($y_i = 0$), $i = 1, \dots, 122$. We applied matrix lasso, lasso, matrix power and power to the data, and evaluated each solution via cross-validation based misclassification error. More specifically, we divided the full data into a training and a testing sample using k -fold cross-validation. When k equals the sample size, it is the leave-one-out cross-validation. Then for the training data, we further employed a 5-fold cross-validation to tune the shrinkage parameter λ . We then applied the tuned model that is fully based on the training data now to the testing data and evaluated the misclassification error rate for testing. Table 5 reports the results for leave-one-out, 5, 10, and 20-fold cross-validation results. It is clearly seen that the matrix version of the regularized estimators achieve smaller misclassification errors compared to the vector version counterparts. Comparing the lasso and the power penalties, the lasso achieves a slightly better classification accuracy.

We also make a few remarks regarding to this data analysis. First, Hung and Wang (2011)

Table 5: Misclassification error rate for EEG data.

Method	leave-one-out	5-fold	10-fold	20-fold
m.lasso	0.230	0.214	0.222	0.181
lasso	0.246	0.287	0.271	0.264
m.power	0.246	0.222	0.228	0.213
power	0.254	0.329	0.244	0.276

analyzed the same data using a matrix logistic regression, which is a rank-1 special case of the tensor GLM of Zhou et al. (2012). Their method does not work for $n < p$, and thus a ridge type regularization was employed. The ridge tuning parameter was chosen so that the leave-one-out classification accuracy is maximized, and the corresponding best classification result was reported, which is slightly better than the classification accuracy reported in Table 5. We believe, however, a fair evaluation of prediction should have the parameter tuning solely based on the training data, and then report the corresponding testing error. If one tunes the parameter based on the reported testing error, the result would be over optimistic. Second, both Li et al. (2010) and Hung and Wang (2011) preprocessed the EEG data by reducing the dimensionality of the matrices to a smaller scale. Both used principal components analysis (PCA) for reduction, but they employed different variants of PCA for matrices. Part of reason for the dimension reduction preprocessing was that their proposed numerical methods can not directly handle the size of 256×64 of the EEG data. By contrast, our proposal is not limited by the matrix size and were directly applied to the original data, since our Nesterov algorithm is highly efficient and scalable. On the other hand, we agree that such preprocessing could potentially improve the overall classification accuracy by removing noisy irrelevant information. However, we note that, PCA is known as an unsupervised dimension reduction approach that can not guarantee full preservation of all relevant information. Moreover, it introduces another layer of tuning, including choosing the right version of PCA for matrices, and determining the optimal number of principal components. Our goal of this data analysis has primarily been the comparison of the regularized matrix regression estimators with the vector counterparts, so we have chosen not to include any preprocessing that requires a significant amount of additional tuning.

6 Discussion

Motivated by modern scientific data arising in areas such as imaging and cytometry, we study in this article the problem of regressions with matrix covariates. Regularization is bound to play a crucial role in such regressions due to the ultrahigh dimensionality and complex structure of the matrix data. We have proposed a class of regularized matrix regression models by penalizing the spectrum of the matrix parameters. It is based upon the observation that the matrix signals are often of, or can be well approximated by, a low rank structure. Consequently, the new method focuses on the sparsity in terms of the rank of the matrix parameters rather than the number of nonzero entries, and is intrinsically different from the classical lasso and related penalization approaches.

In many applications, sparsity is sought in certain pre-specified *basis* systems rather than the original coordinates. Specifically, the systematic component takes the form

$$\eta = \boldsymbol{\gamma}^\top \mathbf{Z} + \langle \mathbf{B}, \mathbf{S}_1^\top \mathbf{X} \mathbf{S}_2 \rangle,$$

where $\mathbf{S}_j \in \mathbb{R}^{p_j \times q_j}$ contains q_j basis vectors for $j = 1, 2$. For 2D images, over-complete wavelet basis ($q_j > p_j$) are often used in both dimensions. For the EEG data that are channels by time points, wavelet or Fourier basis can be applied to the time dimension. It is of direct interest to seek a sparse low rank representation of the signal \mathbf{B} in the basis system by solving the regularized regression. In that sense, our proposed regularized matrix regression can be considered as the matrix analog of the classical basis pursuit problem (Chen et al., 2001).

We have concentrated on problems with matrix covariates throughout this article. In applications such as anatomical magnetic resonance imaging (MRI) and functional magnetic resonance imaging (fMRI), the covariates are in the form of multidimensional arrays, a.k.a. tensors ($D > 2$). It is natural to extend the work here to regularized tensor regressions. However, the problem formulation requires an appropriate norm for tensors that is in analogous to the nuclear norm for matrices, and the involved regularization and optimization are to be fundamentally different from the methods for matrices. We are currently pursuing this line of extension, and will report the results elsewhere.

References

- Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723.
- Armagan, A., Dunson, D., and Lee, J. (2011). Generalized double Pareto shrinkage. *arxiv:1104.0861v1*.
- Beck, A. and Teboulle, M. (2009a). Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems. *Trans. Img. Proc.*, 18(11):2419–2434.
- Beck, A. and Teboulle, M. (2009b). A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sci.*, 2(1):183–202.
- Borwein, J. M. and Lewis, A. S. (2006). *Convex Analysis and Nonlinear Optimization*. CMS Books in Mathematics/Ouvrages de Mathématiques de la SMC, 3. Springer, New York, second edition. Theory and examples.
- Cai, J.-F., Candès, E. J., and Shen, Z. (2010). A singular value thresholding algorithm for matrix completion. *SIAM J. Optim.*, 20(4):1956–1982.
- Candès, E. J. and Recht, B. (2009). Exact matrix completion via convex optimization. *Found. Comput. Math.*, 9(6):717–772.
- Candès, E. J., Wakin, M. B., and Boyd, S. P. (2008). Enhancing sparsity by reweighted l_1 minimization. *J. Fourier Anal. Appl.*, 14(5-6):877–905.
- Chen, S. S., Donoho, D. L., and Saunders, M. A. (2001). Atomic decomposition by basis pursuit. *SIAM Rev.*, 43(1):129–159.
- Chiquet, J., Grandvalet, Y., and Charbonnier, C. (2011). Sparsity with sign-coherent groups of variables via the cooperative-lasso. *Annals of Applied Statistics (to appear)*, arXiv:1103.2697.
- Donoho, D. L. and Johnstone, I. M. (1994). Ideal spatial adaptation by wavelet shrinkage. *Biometrika*, 81(3):425–455.
- Efron, B. (2004). The estimation of prediction error: covariance penalties and cross-validation. *J. Amer. Statist. Assoc.*, 99(467):619–642. With comments and a rejoinder by the author.
- Efron, B., Hastie, T., Johnstone, I., and Tibshirani, R. (2004). Least angle regression. *Ann. Statist.*, 32(2):407–499. With discussion, and a rejoinder by the authors.
- Fan, J. and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.*, 96(456):1348–1360.
- Frank, I. E. and Friedman, J. H. (1993). A statistical view of some chemometrics regression tools. *Technometrics*, 35(2):109–135.

- Golub, G. H. and Van Loan, C. F. (1996). *Matrix Computations*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, third edition.
- Hoerl, A. E. and Kennard, R. W. (1970). Ridge regression: biased estimation for nonorthogonal problems. *Technometrics*, 12:55–67.
- Hung, H. and Wang, C.-C. (2011). Matrix variate logistic regression analysis. *arXiv:1105.2150v1*.
- Johnstone, I. M. (2001). On the distribution of the largest eigenvalue in principal components analysis. *Ann. Statist.*, 29(2):295–327.
- Lange, K. (2004). *Optimization*. Springer Texts in Statistics. Springer-Verlag, New York.
- Li, B., Kim, M. K., and Altman, N. (2010). On dimension folding of matrix- or array-valued statistical objects. *The Annals of Statistics*, 38(2):1094–1121.
- Mazumder, R., Hastie, T., and Tibshirani, R. (2010). Spectral regularization algorithms for learning large incomplete matrices. *Journal of Machine Learning Research*, 11:2287–2322.
- McCullagh, P. and Nelder, J. A. (1983). *Generalized Linear Models*. Monographs on Statistics and Applied Probability. Chapman & Hall, London.
- Nemirovski, A. (1994). Efficient methods in convex programming.
- Nesterov, Y. (1983). A method of solving a convex programming problem with convergence rate $O(1/k^2)$. *Soviet Mathematics Doklady*, 27(2):372–376.
- Nesterov, Y. (2004). *Introductory Lectures on Convex Optimization*, volume 87 of *Applied Optimization*. Kluwer Academic Publishers, Boston, MA. A basic course.
- Nocedal, J. and Wright, S. J. (2006). *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer, New York, second edition.
- Qiu, P. (2005). *Image Processing and Jump Regression Analysis*. New York: John Wiley & Sons.
- Qiu, P. (2007). Jump surface estimation, edge detection, and image restoration. *Journal of American Statistical Association*, 102(745-756).
- Recht, B., Fazel, M., and Parrilo, P. A. (2010). Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Rev.*, 52(3):471–501.
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464.
- Stein, C. M. (1981). Estimation of the mean of a multivariate normal distribution. *Ann. Statist.*, 9(6):1135–1151.

- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B*, 58(1):267–288.
- Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., and Knight, K. (2005). Sparsity and smoothness via the fused lasso. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 67(1):91–108.
- Yuan, M. and Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 68(1):49–67.
- Zhang, C.-H. (2010). Nearly unbiased variable selection under minimax concave penalty. *Ann. Statist.*, 38(2):894–942.
- Zhou, H., Li, L., and Zhu, H. (2012). Tensor regression with applications in neuroimaging data analysis. *arXiv:1203.3209*.
- Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 67(2):301–320.
- Zou, H., Hastie, T., and Tibshirani, R. (2007). On the “degrees of freedom” of the lasso. *Ann. Statist.*, 35(5):2173–2192.